arXiv:1605.08108v1 [math.OC] 26 May 2016

FLAG: Fast Linearly-Coupled Adaptive Gradient Method

Xiang Cheng[*1], Farbod Roosta-Khorasani[†2,3], Peter L. Bartlett[‡1,2] and Michael W. Mahoney[§2,3]

[1]Department of Computer Science, UC Berkeley, Berkeley, CA
[2]Department of Statistics, UC Berkeley, Berkeley, CA
[3]International Computer Science Institute, Berkeley, CA

May 27, 2016

Abstract

The celebrated Nesterov's accelerated gradient method offers great speed-ups compared to the classical gradient descend method as it attains the optimal first-order oracle complexity for smooth convex optimization. On the other hand, the popular AdaGrad algorithm competes with mirror descent under the best regularizer by adaptively scaling the gradient. Recently, it has been shown that the accelerated gradient descent can be viewed as a linear combination of gradient descent and mirror descent steps. Here, we draw upon these ideas and present a fast linearly-coupled adaptive gradient method (FLAG) as an accelerated version of AdaGrad, and show that our algorithm can indeed offer the best of both worlds. Like Nesterov's accelerated algorithm and its proximal variant, FISTA, our method has a convergence rate of $1/T^2$ after T iterations. Like AdaGrad our method adaptively chooses a regularizer, in a way that performs almost as well as the best choice of regularizer in hindsight.

1 Introduction

We consider the problem of minimizing a convex function F over a convex set $\mathcal{C} \subset \mathbb{R}^d$, and we focus on first order methods, which exploit only value and gradient information about

[*]x.cheng@berkeley.edu
[†]farbod@stat.berkeley.edu
[‡]bartlett@cs.berkeley.edu
[§]mmahoney@stat.berkeley.edu

F. These methods have become very important for many of the large-scale optimization problems that arise in machine learning applications. Two techniques have emerged as powerful tools for large-scale optimization: First, Nesterov's accelerated algorithms [11] and its proximal variants, Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [2, 3], can exploit smoothness to improve the convergence rate of a simple gradient/ISTA method from accuracy $\mathcal{O}(1/T)$ to $\mathcal{O}(1/T^2)$ after T iterations. Second, adaptive regularization approaches such as AdaGrad [7] can optimize a gradient method's step-size in different directions, in some sense making the optimization problem better-conditioned. These two techniques have also become popular in non-convex optimization problems, such as parameter estimation in deep neural networks (see, for example,the work of [13] on the benefits of acceleration, and AdaGrad [7, 6], RMSProp [14], ESGD [5], and Adam [9] on adaptive regularization methods in this context).

In this paper, we introduce a new algorithm, called *FLAG*, that combines the benefits of these two techniques. Like Nesterov's accelerated algorithms, our method has a convergence rate of $1/T^2$. Like AdaGrad, our method adaptively chooses a regularizer, in a way that performs almost as well as the best choice of regularizer in hindsight. In addition, our improvement over FISTA is roughly a square of AdaGrad's improvement over mirror descent (see section 2.1 for details).

There have been a number of papers in recent years dealing with ways to interpret acceleration. Our algorithm is heavily influenced by [1]. This insightful paper introduced a view of accelerated gradient descent as a linear combination of gradient and mirror descent steps. We exploit this viewpoint in introducing adaptive regularization to our accelerated method. Authors in [1] use a fixed schedule to combine gradient and mirror descent steps. However, [4] views acceleration as an ellipsoid-like algorithm, and instead, determines an appropriate ratio using a line search. This latter idea is also crucial for our algorithm, as we do not know the right ratio in advance.

The optimal stepsize for an accelerated method generally depends on both the smoothness parameter of the objective function as well as the properties of the regularizer. In FLAG, the regularizer is chosen adaptively, and we show that it is also possible to adapt the stepsize in a way that is competitive with the best choice in hindsight. In fact, our method for picking the adaptive stepsizes is closely related to the approach in [8]. There, the authors considered the problem of picking the right stepsize to adapt to an unknown strong-convexity parameter. We use a related approach, but with different proof techniques, to choose the stepsize to adapt to a changing smoothness parameter.

Finally, it should also be noted that there are some interesting papers that study the continuous-time limit of acceleration algorithms, e.g., [12, 10, 15]. Indeed, studying adaptive regularization in the continuous time setting is an interesting future direction for research.

1.1 Notation and Definitions

In what follows, vectors are considered as column vectors and are denoted by bold lower case letters, e.g., $\mathbf{x}$ and matrices are denoted by regular capital letters, e.g., A. We overload

the "diag" operator as follows: for a given matrix A and a vector $\mathbf{x}$, $\mathrm{diag}(A)$ and $\mathrm{diag}(\mathbf{x})$ denote the vector made from the diagonal elements of A and a diagonal matrix made from the elements of $\mathbf{v}$, respectively. Vector norms $\|\mathbf{x}\|_1$ $\|\mathbf{x}\|_2$ and $\|\mathbf{x}\|_\infty$ denote the standard ℓ_1, ℓ_2 and ℓ_∞ respectively. We adopt the `Matlab` notation for accessing the elements of vectors and matrices, i.e., i^{th} components of a vector $\mathbf{x}$ is indicated by $\mathbf{x}(i)$ and $A(i,:)$ denotes the entire i^{th} row of the matrix A. Finally, $A_k \leftarrow [A_{k-1}, \mathbf{v}]$ signifies that A_k is the augmentation of the matrix A_{k-1} with the column vector $\mathbf{v}$.

Consider the optimization problem

$$\min_{\mathbf{x}\in\mathcal{C}} F(\mathbf{x}) = f(\mathbf{x}) + h(\mathbf{x}), \tag{1}$$

where $f : \mathbb{R}^d \to \mathbb{R}$ and $h : \mathbb{R}^p \to \mathbb{R}$ are closed proper convex convex functions and $\mathcal{C}$ is a closed convex set. We further assume that f is differentiable with L-Lipschitz gradient, i.e.,

$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\|_2 \leq L\|\mathbf{x} - \mathbf{y}\|_2, \quad \forall \mathbf{x}, \mathbf{y} \in \mathcal{C}, \tag{2}$$

and we allow h to be (possibly) non-smooth with sub-differential at $\mathbf{x}$ denoted by $\partial h(\mathbf{x})$.

The proximal operator [11] associated with f, h and L is defined as

$$\mathbf{prox}(\mathbf{x}) := \arg\min_{\mathbf{y}\in\mathcal{C}} h(\mathbf{y}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x}\rangle + \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|_2^2 \tag{3a}$$

$$= \arg\min_{\mathbf{y}\in\mathcal{C}} \; h(\mathbf{y}) + \frac{L}{2}\|\mathbf{y} - \big(\mathbf{x} - \frac{1}{L}\nabla f(\mathbf{x})\big)\|_2^2. \tag{3b}$$

For a symmetric positive definite(SPD) matrix S, define $\psi(\mathbf{x}) := \frac{1}{2}\mathbf{x}^T S\mathbf{x}$. Note that ψ is 1-strongly convex with respect to the norm $\|\mathbf{x}\|_S := \sqrt{\mathbf{x}^T S\mathbf{x}}$, i.e., $\forall \mathbf{x}, \mathbf{y} \in \mathcal{C}$,

$$\psi(\mathbf{x}) \geq \psi(\mathbf{y}) + \langle \nabla\psi(\mathbf{y}), \mathbf{x} - \mathbf{y}\rangle + \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|_S^2.$$

The Bregman divergence associated with ψ, is define as

$$B_\psi(\mathbf{x}, \mathbf{y}) := \psi(\mathbf{x}) - \psi(\mathbf{y}) - \langle \nabla\psi(\mathbf{y}), \mathbf{x} - \mathbf{y}\rangle = \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|_S^2.$$

We will use the fact that, for any $\mathbf{z} \in \mathcal{C}$,

$$\langle S(\mathbf{x} - \mathbf{y}), \mathbf{x} - \mathbf{z}\rangle = \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|_S^2 + \frac{1}{2}\|\mathbf{x} - \mathbf{z}\|_S^2 - \frac{1}{2}\|\mathbf{y} - \mathbf{z}\|_S^2. \tag{4}$$

It is easy to see that the dual of $\psi(\mathbf{x})$ is given by

$$\psi^*(\mathbf{x}) := \sup_{\mathbf{v}\in\mathbb{R}^d} \langle \mathbf{x}, \mathbf{v}\rangle - \psi(\mathbf{v}) = \frac{1}{2}\mathbf{x}^T S^{-1}\mathbf{x}. \tag{5}$$

2 The Algorithm and Main Result

In this section, we describe our main algorithm, FLAG, and give its convergence properties in Theorem 1. The core of FLAG consists of five essential ingredients:

(i) A proximal gradient step (Step 3)
(ii) Construction of the adaptive regularization (Steps 4-9)
(iii) Picking the Adaptive stepsize (Step 11)
(iv) A mirror descent step (Step 12)
(v) Linear combination of the proximal gradient step and the mirror descent step (Step 13)

Algorithm 1 FLAG

INPUT: $\mathbf{x}_1 = \mathbf{y}_1 = \mathbf{z}_1$, $\eta_0 = 0$, $G_0 = [\text{ empty }]$, T

1: $\epsilon \leftarrow 1/\left(6dT^3\right)$
2: **for** k=1 to T **do**
3: $\quad \mathbf{y}_{k+1} \leftarrow \mathbf{prox}(\mathbf{x}_k)$
4: $\quad \mathbf{p}_k \leftarrow -L(\mathbf{y}_{k+1} - \mathbf{x}_k)$
5: $\quad \mathbf{g}_k \leftarrow \mathbf{p}_k / \|\mathbf{p}_k\|_2$
6: $\quad G_k \leftarrow \left[G_{k-1}, \mathbf{g}_k\right]$
7: $\quad \mathbf{s}_k(i) \leftarrow \|G_k(i,:)\|_2$
8: $\quad S_k \leftarrow \mathrm{diag}(\mathbf{s}_k) + \delta\mathbb{I}$
9: $\quad \psi_k(\mathbf{x}) \leftarrow \frac{1}{2}\mathbf{x}^T S_k \mathbf{x}$
10: $\quad L_k \leftarrow L\mathbf{g}_k^T S_k^{-1} \mathbf{g}_k$
11: $\quad \eta_k \leftarrow \frac{1}{2L_k} + \sqrt{\frac{1}{4L_k^2} + \frac{\eta_{k-1}^2 L_{k-1}}{L_k}}$
12: $\quad \mathbf{z}_{k+1} \leftarrow \arg\min_{\mathbf{z}\in\mathcal{C}} \langle \eta_k \mathbf{p}_k, \mathbf{z} - \mathbf{z}_k\rangle + \frac{1}{2}\|\mathbf{z} - \mathbf{z}_k\|_{S_k}^2$
13: $\quad \mathbf{x}_{k+1} \leftarrow \mathrm{BinarySearch}(\mathbf{z}_{k+1}, \mathbf{y}_{k+1}, \epsilon)$
return $\mathbf{y}_{T+1}$

The subroutine *BinarySearch* is given in Algorithm 2, where $Bisection(r, 0, 1, \epsilon)$ is the usual bisection routine for finding the root of a single variable function $r(t)$ in the interval $(0, 1)$ and to the accuracy of ϵ. More specifically, for a root r^* such that $r(t^*) = 0$ and given $r(0)r(1) < 0$, the sub-routine $Bisection(r, 0, 1, \epsilon)$ returns an approximation $t \in (0, 1)$ to t^* such that $|t - t^*| \le \epsilon$ and this is done with only $\log(1/\epsilon)$ function evaluations.

Algorithm 2 BinarySearch

INPUT: $\mathbf{z}$, $\mathbf{y}$, and ϵ

1: Define the univariate function $r(t) := \langle \mathbf{prox}\left(t\mathbf{y} + (1-t)\mathbf{z}\right) - \left(t\mathbf{y} + (1-t)\mathbf{z}\right), \mathbf{y} - \mathbf{z}\rangle$
2: **if** $r(1) \ge 0$ **then return** $\mathbf{y}$
3: **if** $r(0) \le 0$ **then return** $\mathbf{z}$
4: $t = \mathrm{Bisection}(r, 0, 1, \epsilon)$
5: **return** $\mathbf{x} = t\mathbf{y} + (1-t)\mathbf{z}$

Theorem 1, gives the convergence properties of Algorithm 1:

Theorem 1 (Main Result: Convergence Property of FLAG)
Let $D := \sup_{\mathbf{x},\mathbf{y}\in\mathcal{C}} \|\mathbf{x}-\mathbf{y}\|_\infty^2$. *For any* $\mathbf{u}\in\mathcal{C}$, *after* T *iterations of Algorithm 1, we get*

$$F(\mathbf{y}_{T+1}) - F(\mathbf{u}) \leq \mathcal{O}\left(\frac{q_T^2}{T}\frac{LD}{T^2}\right),$$

where $q_T \in [\sqrt{T}, \sqrt{Td}]$. *Also each iteration takes* $\mathcal{O}(d\log(dT^3))$ *time.*

2.1 Comparison of FLAG with AdaGrad and FISTA

In this Section, we briefly describe the advantages of FLAG relative to AdaGrad and FISTA.

Let ∇_k be the subgradient of F at iteration k, $\mathcal{V} := \{S \in \mathbb{R}^{d\times d} \mid S \text{ is diagonal, } S_{ii} > 0, \text{ trace}(S) \leq d\}$, $D_2 := \sup_{\mathbf{x},\mathbf{y}\in\mathcal{C}} \|\mathbf{x}-\mathbf{y}\|_2^2$, and $J_A := \inf_{S\in\mathcal{V}} \sum_{k=1}^T \nabla_k^T S^{-1}\nabla_k / \sum_{k=1}^T \|\nabla_k^T\|_2^2$. Using a similar argument to that used in Lemma 5, we see that $J_A \in [1/d, 1]$. AdaGrad achieves a rate of

$$F(\bar{\mathbf{y}}) - F(\mathbf{x}^*) \leq O\left(\sqrt{J_A}\frac{\sqrt{Dd\sum_{k=1}^T \|g_k^T\|_2^2}}{T}\right), \tag{6}$$

which is in contrast to mirror descent, which achieves

$$F(\bar{\mathbf{y}}) - F(\mathbf{x}^*) \leq O\left(\frac{\sqrt{D_2\sum_{k=1}^T \|g_k^T\|_2^2}}{T}\right),$$

where $\bar{\mathbf{y}}$ is the average of all $\mathbf{y}_k$'s. Thus, one can see that the improvement of AdaGrad compared to mirror descent is $\sqrt{J_A Dd/D_2}$.

Now, let $\mathbf{g}_k$ be defined as in Algorithm 1. One can verify using Lemma 5 that $J_B := q_T^2/T = \inf_{S\in\mathcal{V}} \sum_{k=1}^T \mathbf{g}_k^T S^{-1}\mathbf{g}_k / \sum_{k=1}^T \|\mathbf{g}_k^T\|_2^2$ and $J_B \in [1/d, 1]$ (J_B is analogous to J_A). Thus, from Theorem 1, we get

$$F(\mathbf{y}_{T+1}) - F(\mathbf{x}^*) \leq O\left(J_B\frac{LDd}{T^2}\right). \tag{7}$$

This is in contrast to FISTA, which achieves

$$F(\mathbf{y}_{T+1}) - F(\mathbf{x}^*) \leq O\left(\frac{LD_2 d}{T^2}\right).$$

As a result, the improvement of FLAG compared to FISTA is $J_B Dd/D_2$. FLAG's improvement over FISTA can thus be up to a square of AdaGrad's improvement over mirror descent. Here, we note that J_A and J_B though analogous, are not the same.

Finally, we can directly compare the rates of (6) and (7) to see the speed-up offered by FLAG over AdaGrad. In particular, FLAG enjoys from the optimal rate of $\mathcal{O}(1/T^2)$, compared to the sub-optimal rate of $\mathcal{O}(1/\sqrt{T})$ of AdaGrad. However, we do stress that AdaGrad does not make any smoothness assumptions, and thus works in more general settings than FLAG.

3 Proofs

Below we give the details for the proof of our main result. The proofs of technical lemmas in Section 3.1 are given in Appendix A.

The proof sketch of the convergence of FLAG is as follows.

(I) FLAG is essentially a combination of mirror descent and proximal gradient descent steps (see Lemmas 1 and 4).

(II) L_k in Algorithm 1 plays the role of an "effective gradient lipschitz constant" in each iteration. The convergence rate of our algorithm ultimately depends on $\sum_{k=1}^{T} L_k = L\sum_{k=1}^{T} \mathbf{g}_k^T S_k^{-1} \mathbf{g}_k$. (see Lemmas 7 and 2)

(III) By picking S_k adaptively, as in AdaGrad, we achieve a non-trivial upper bound for $\sum_{k=1}^{T} L_k$. (see Lemma 5)

(IV) Our algorithm relies on picking an $\mathbf{x}_k$ at each iteration that satisfies an inequality involving L_k (see Corollary 1). However, because L_k is not known at the moment of picking $\mathbf{x}_k$, we must choose an $\mathbf{x}_k$ to roughly satisfy the inequality for all possible values of L_k. We do this by picking $\mathbf{x}_k$ using binary search. (see Lemmas 2 and 3 and Corollary 1)

(V) Finally, we need to pick the right stepsizes for each iteration. Our scheme is very similar to the one used in [1], but generalized to handle a different L_k each iteration. (see Lemma 6 and 7 and Corollary 2).

Theorem 2 combines items I, II, and IV. Theorem 1 combines Theorem 2 with items III and V to get our final result.

3.1 Technical Lemmas

We have the following key result (see Lemma 2.3 [2]) regarding the vector $\mathbf{p} = -L(\mathbf{prox}(\mathbf{x}) - \mathbf{x})$, as in Step 4 of FLAG, which is called the *Gradient Mapping* of F on $\mathcal{C}$.

Lemma 1 (Gradient Mapping)
For any $\mathbf{x}, \mathbf{y} \in \mathcal{C}$*, we have*

$$F(\boldsymbol{prox}(\mathbf{x})) \leq F(\mathbf{y}) + \langle L(\boldsymbol{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{x}\rangle - \frac{L}{2}\|\mathbf{x} - \boldsymbol{prox}(\mathbf{x})\|_2^2,$$

where $\boldsymbol{prox}(\mathbf{x})$ *is defined as in* (3)*. In particular,* $F(\boldsymbol{prox}(\mathbf{x})) \leq F(\mathbf{x}) - \frac{L}{2}\|\mathbf{x} - \boldsymbol{prox}(\mathbf{x})\|_2^2$.

The following lemma establishes the Lipschitz continuity of the **prox** operator.

Lemma 2 (Prox Operator Continuity)
$\boldsymbol{prox} : \mathbb{R}^d \to \mathbb{R}^d$ *is a* 2*-Lipschitz continuous, that is, for any* $\mathbf{x}, \mathbf{y} \in \mathcal{C}$*, we have*

$$\|\boldsymbol{prox}(\mathbf{x}) - \boldsymbol{prox}(\mathbf{y})\|_2 \leq 2\|\mathbf{x} - \mathbf{y}\|_2.$$

Using **prox** operator continuity Lemma 2, we can conclude that given any $\mathbf{y}, \mathbf{z} \in \mathcal{C}$, if $\langle \mathbf{prox}(\mathbf{y}) - \mathbf{y}, \mathbf{y} - \mathbf{z}\rangle < 0$ and $\langle \mathbf{prox}(\mathbf{z}) - \mathbf{z}, \mathbf{y} - \mathbf{z}\rangle > 0$, then there must be a $t^* \in (0, 1)$ for which $\mathbf{w} = t^*\mathbf{y} + (1 - t^*)\mathbf{z}$ gives $\langle \mathbf{prox}(\mathbf{w}) - \mathbf{w}, \mathbf{y} - \mathbf{z}\rangle = 0$. Algorithm 2 finds an approximation to $\mathbf{w}$ in $\mathcal{O}(\log L/\epsilon)$ iterations.

Lemma 3 (Binary Search Lemma)
Let $\mathbf{x} = \mathit{BinarySearch}(\mathbf{z}, \mathbf{y}, \epsilon)$ *defined as in Algorithm 2. Then one of 3 cases happen:*
(i) $\mathbf{x} = \mathbf{y}$ *and* $\langle \boldsymbol{prox}(\mathbf{x}) - \mathbf{x}, \mathbf{x} - \mathbf{z}\rangle \geq 0$*,*
(ii) $\mathbf{x} = \mathbf{z}$ *and* $\langle \boldsymbol{prox}(\mathbf{x}) - \mathbf{x}, \mathbf{y} - \mathbf{x}\rangle \leq 0$*, or*
(iii) $\mathbf{x} = t\mathbf{y} + (1 - t)\mathbf{z}$ *for some* $t \in (0, 1)$ *and* $|\langle \boldsymbol{prox}(\mathbf{x}) - \mathbf{x}, \mathbf{y} - \mathbf{z}\rangle| \leq 3\|\mathbf{y} - \mathbf{z}\|_2^2\epsilon$.

Using the above result, we can prove the following:

Corollary 1
Let $\mathbf{x}_k$*,* $\mathbf{y}_k$*,* $\mathbf{z}_k$ *and* ϵ_k *be defined as in Algorithm 1 and* $\eta_k L_k \geq 1$*. Then for all* $k \geq 1$*,*

$$\langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k\rangle \leq (\eta_k L_k - 1)\langle \mathbf{p}_k, \mathbf{y}_k - \mathbf{x}_k\rangle + \frac{DL\eta_k L_k}{T^3}.$$

Next, we state a result regarding the mirror descent step. Similar results can be found in most texts on online optimization, e.g. [1].

Lemma 4 (Mirror Descent Inequality)
Let $\mathbf{z}_{k+1} = \arg\min_{\mathbf{z}\in\mathcal{C}} \langle \eta_k \mathbf{p}_k, \mathbf{z} - \mathbf{z}_k\rangle + \frac{1}{2}\|\mathbf{z} - \mathbf{z}_k\|_{S_k}^2$ *and* $D := \sup_{\mathbf{x},\mathbf{y}\in\mathcal{C}} \|\mathbf{x} - \mathbf{y}\|_\infty^2$ *be the diameter of* $\mathcal{C}$ *measured by infinity norm. Then for any* $\mathbf{u} \in \mathcal{C}$*, we have*

$$\sum_{k=1}^{T} \langle \eta_k \mathbf{p}_k, \mathbf{z}_k - u\rangle \leq \sum_{k=1}^{T} \frac{\eta_k^2}{2} \|\mathbf{p}_k\|_{S_k^*}^2 + \frac{D}{2}\|\mathbf{s}_T\|_1$$

Finally, we state a similar result to that of [7] that captures the benefits of using S_k in FLAG.

Lemma 5 (AdaGrad Inequalities)
Define $q_T := \sum_{i=1}^{d} \|G_T(i,:)\|_2$*, where* G_k *is as in Step 6 of Algorithm 1. We have*
(i) $\sum_{k=1}^{T} \mathbf{g}_k^T S_k^{-1} \mathbf{g}_k \leq 2q_T$,
(ii) $q_T^2 = \min_{S\in\mathcal{S}} \sum_{k=1}^{T} \mathbf{g}_k^T S^{-1} \mathbf{g}_k$*, where* $\mathcal{S} := \{S \in \mathbb{R}^{d\times d} \mid S \text{ is diagonal}, S_{ii} > 0, \text{trace}(S) \leq 1\}$*, and*
(iii) $\sqrt{T} \leq q_T \leq \sqrt{dT}$.

3.2 Master Theorem

We can now prove the central theorems of this section.

Theorem 2
Let $D := \sup_{\mathbf{x},\mathbf{y}\in\mathcal{C}} \|\mathbf{x} - \mathbf{y}\|_\infty^2$*. For any* $\mathbf{u} \in \mathcal{C}$*, after* T *iterations of Algorithm 1, we get*

$$\sum_{k=1}^{T} \Big\{ \left(\eta_{k-1}^2 L_{k-1} - \eta_k^2 L_k + \eta_k\right) F(\mathbf{y}_k) - \eta_k F(\mathbf{u}) \Big\} + \eta_T^2 L_T F(\mathbf{y}_{T+1}) \leq \sum_{k=1}^{T} \frac{DL\eta_k^2 L_k}{T^3} + \frac{D}{2}\|\mathbf{s}_T\|_1.$$

Proof of Theorem 2 Noting that $\mathbf{p}_k = -L(\mathbf{y}_{k+1} - \mathbf{x}_k)$ is the gradient mapping of F on $\mathcal{C}$, it follows that

$$\begin{aligned}
\sum_{k=1}^{T} \eta_k (F(\mathbf{y}_{k+1}) - F(\mathbf{u})) &= \sum_{k=1}^{T} \eta_k (F(\mathbf{prox}(\mathbf{x}_k)) - F(\mathbf{u})) \\
\text{(By the Gradient Mapping Lemma 1)} \quad &\leq \sum_{k=1}^{T} \eta_k \langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{u} \rangle - \frac{\eta_k}{2L} \|\mathbf{p}_k\|_2^2 \\
&= \sum_{k=1}^{T} \eta_k \langle \mathbf{p}_k, (\mathbf{z}_k - \mathbf{u}) \rangle + \sum_{k=1}^{T} \eta_k \langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k \rangle - \frac{\eta_k}{2L} \|\mathbf{p}_k\|_2^2 \\
\text{(by Lemma 4)} \quad &\leq \sum_{k=1}^{T} \frac{\eta_k^2}{2} \|\mathbf{p}_k\|_{S_k^*}^2 + \frac{D}{2} \|\mathbf{s}_T\|_1 \\
&\qquad + \sum_{k=1}^{T} \eta_k \langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k \rangle - \frac{\eta_k}{2L} \|\mathbf{p}_k\|_2^2 \\
\text{(By Step 10 of Algorithm 1)} \quad &= \sum_{k=1}^{T} \frac{\eta_k (\eta_k L_k - 1)}{2L} \|\mathbf{p}_k\|_2^2 + \frac{D}{2} \|\mathbf{s}_T\|_1 + \sum_{k=1}^{T} \eta_k \langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k \rangle \\
\text{(By Corollary 1)} \quad &\leq \sum_{k=1}^{T} \frac{\eta_k (\eta_k L_k - 1)}{2L} \|\mathbf{p}_k\|_2^2 + \frac{D}{2} \|\mathbf{s}_T\|_1 \\
&\qquad + \sum_{k=1}^{T} \left(\eta_k (\eta_k L_k - 1) \langle \mathbf{p}_k, \mathbf{y}_k - \mathbf{x}_k \rangle + \frac{DL\eta_k^2 L_k}{T^3} \right) \\
\text{(By the Gradient Mapping Lemma 1)} \quad &\leq \sum_{k=1}^{T} \frac{DL\eta_k^2 L_k}{T^3} + \frac{D}{2} \|\mathbf{s}_T\|_1 \\
&\qquad + \sum_{k=1}^{T} \eta_k (\eta_k L_k - 1) \left(F(\mathbf{y}_k) - F(\mathbf{y}_{k+1}) \right).
\end{aligned}$$

Now rearranging terms and re-indexing the summations gives the desired result. ∎

3.3 Choosing the Stepsize

In this section, we discuss the final piece of our algorithm: choosing the stepsize η_k for the mirror descent step.

Lemma 6
For the choice of η_k in Algorithm 1 and $k \geq 1$, we have

(i) $\eta_k^2 L_k = \sum_{i=1}^k \eta_i$,

(ii) $\eta_{k-1}^2 L_{k-1} - \eta_k^2 L_k + \eta_k = 0$, *and*

(iii) $\eta_k L_k \geq 1$.

Proof We prove (i) by induction. For $k = 1$, is is easy to verify that $\eta_1 = 1/L_1$, and so $\eta_1^2 L_1 = \eta_1$ and the base case follows trivially. Now suppose $\eta_{k-1}^2 L_{k-1} = \sum_{i=1}^{k-1} \eta_i$. Rearranging (i) for k gives

$$0 = \eta_k^2 L_k - \eta_k - \sum_{i=1}^{k-1} \eta_i = \eta_k^2 L_k - \eta_k - \eta_{k-1}^2 L_{k-1}.$$

Now, it is easy to verify that the choice of η_k in Algorithm 1 is a solution of the above quadratic equation. The rest of the items follow immediately from part i. ∎

Corollary 2
Let $D := \sup_{\mathbf{x},\mathbf{y} \in \mathcal{C}} \|\mathbf{x} - \mathbf{y}\|_\infty^2$. *For any* $\mathbf{u} \in \mathcal{C}$, *after* T *iterations of Algorithm 1, we get*

$$F(\mathbf{y}_{T+1}) - F(\mathbf{u}) \leq \frac{LD}{T^2} + \frac{D\|\mathbf{s}_T\|_1}{2\sum_{k=1}^T \eta_k}.$$

Proof of corollary 2 Combining Theorem 2 and Lemma 6, and noting that $\eta_k^2 L_k = \sum_{i=1}^k \eta_i \leq \sum_{i=1}^T \eta_i = \eta_T^2 L_T$ gives the desired result. ∎

Finally, it only remains to lower bound $\sum_{k=1}^T \eta_k$, which is done in the following Lemma.

Lemma 7
For the choice of η_k *in Algorithm 1, we have*

$$\sum_{k=1}^T \eta_k \geq \frac{T^3}{1000 \sum_{k=1}^T L_k}$$

Remark We note here that we made little effort to minimize constants, and that we used rather sloppy bounds such as $T-1 \geq T/2$. As a result, the constant appearing above is very conservative and a mere by product of our proof technique. We have numerically verified that a much smaller constant (e.g., 10) indeed satisfies the bound above.

And now, the proof of our main result, Theorem 1, follows rather immediately:

Proof of Theorem 1 The result follows immediately from Lemma 7 and Corollary 2 and noting that $\sum_{k=1}^{T} L_k = L \sum_{k=1}^{T} \mathbf{g}_k^T S_k^{-1} \mathbf{g} \leq 2Lq_{_T}$ by Lemma 5 and $\|\mathbf{s}_T\|_1 = q_{_T}$ by Step 7 of Algorithm 1 and definition of $q_{_T}$ in Lemma 5. This gives

$$F(\mathbf{y}_{T+1}) - F(\mathbf{u}) \leq \frac{LD}{T^2} + \frac{q_{_T}^2}{T} \frac{1000LD}{T^2} \leq \frac{q_{_T}^2}{T} \frac{1001LD}{T^2}$$

The run-time per iteration follows from having to do $\log_2(\frac{1}{\epsilon})$ calls to bisection, each taking $O(n)$ time. ∎

References

[1] Zeyuan Allen-Zhu and Lorenzo Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. *arXiv preprint arXiv:1407.1537*, 2014.

[2] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.

[3] Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

[4] Sébastien Bubeck, Yin Tat Lee, and Mohit Singh. A geometric alternative to Nesterov's accelerated gradient descent. *arXiv preprint arXiv:1506.08187*, 2015.

[5] Yann Dauphin, Harm de Vries, and Yoshua Bengio. Equilibrated adaptive learning rates for non-convex optimization. In *Advances in Neural Information Processing Systems*, pages 1504–1512, 2015.

[6] Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Andrew Senior, Paul Tucker, Ke Yang, Quoc V Le, et al. Large scale distributed deep networks. In *Advances in Neural Information Processing Systems*, pages 1223–1231, 2012.

[7] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.

[8] Elad Hazan, Alexander Rakhlin, and Peter L Bartlett. Adaptive online gradient descent. In *Advances in Neural Information Processing Systems*, pages 65–72, 2007.

[9] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[10] Walid Krichene, Alexandre Bayen, and Peter L Bartlett. Accelerated mirror descent in continuous and discrete time. In *Advances in Neural Information Processing Systems*, pages 2827–2835, 2015.

[11] Neal Parikh and Stephen P Boyd. Proximal algorithms. *Foundations and Trends in optimization*, 1(3):127–239, 2014.

[12] Weijie Su, Stephen Boyd, and Emmanuel Candes. A differential equation for modeling Nesterovs accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*, pages 2510–2518, 2014.

[13] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th international conference on machine learning (ICML-13)*, pages 1139–1147, 2013.

[14] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4, 2012.

[15] Andre Wibisono, Ashia C Wilson, and Michael I Jordan. A variational perspective on accelerated methods in optimization. *arXiv preprint arXiv:1603.04245*, 2016.

A Proofs of Technical Lemmas

Proof of Lemma 1 This result is the same as Lemma 2.3 in [2]. We bring its proof here for completeness.

For any $\mathbf{y} \in \mathcal{C}$, any sub-gradient, $\mathbf{v}$, of h at $\mathbf{prox}(\mathbf{x})$, i.e., $\mathbf{v} \in \partial h(\mathbf{prox}(\mathbf{x}))$, and by optimality of $\mathbf{prox}(\mathbf{x})$ in (3), we have

$$\begin{aligned} 0 &\leq \langle \nabla f(\mathbf{x}) + \mathbf{v} + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{prox}(\mathbf{x}) \rangle \\ &= \langle \nabla f(\mathbf{x}) + \mathbf{v} + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \langle \nabla f(\mathbf{x}) + \mathbf{v} + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{x} - \mathbf{prox}(\mathbf{x}) \rangle, \end{aligned}$$

and so

$$\begin{aligned} \langle \nabla f(\mathbf{x}), \mathbf{prox}(\mathbf{x}) - \mathbf{x} \rangle \leq \langle \nabla f(\mathbf{x}) + \mathbf{v} + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \\ + \langle \mathbf{v}, \mathbf{x} - \mathbf{prox}(\mathbf{x}) \rangle - L\|\mathbf{x} - \mathbf{prox}(\mathbf{x})\|_2^2, \end{aligned}$$

Now from L-Lipschitz continuity of ∇f as well as convexity of f and h, we get

$$\begin{aligned} F(\mathbf{prox}(\mathbf{x})) &= f(\mathbf{prox}(\mathbf{x})) + h(\mathbf{prox}(\mathbf{x})) \\ &\leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{prox}(\mathbf{x}) - \mathbf{x} \rangle + \frac{L}{2}\|\mathbf{prox}(\mathbf{x}) - \mathbf{x}\|_2^2 + h(\mathbf{prox}(\mathbf{x})) \\ &\leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}) + \mathbf{v} + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \langle \mathbf{v}, \mathbf{x} - \mathbf{prox}(\mathbf{x}) \rangle \\ &\qquad - \frac{L}{2}\|\mathbf{x} - \mathbf{prox}(\mathbf{x})\|_2^2 + h(\mathbf{prox}(\mathbf{x})) \\ &\leq f(\mathbf{y}) + \langle \mathbf{v} + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \langle \mathbf{v}, \mathbf{x} - \mathbf{prox}(\mathbf{x}) \rangle \\ &\qquad - \frac{L}{2}\|\mathbf{x} - \mathbf{prox}(\mathbf{x})\|_2^2 + h(\mathbf{prox}(\mathbf{x})) \\ &= f(\mathbf{y}) + \langle L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \langle \mathbf{v}, \mathbf{y} - \mathbf{prox}(\mathbf{x}) \rangle \\ &\qquad - \frac{L}{2}\|\mathbf{x} - \mathbf{prox}(\mathbf{x})\|_2^2 + h(\mathbf{prox}(\mathbf{x})) \\ &\leq F(\mathbf{y}) + \langle L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{y} - \mathbf{x} \rangle - \frac{L}{2}\|\mathbf{x} - \mathbf{prox}(\mathbf{x})\|_2^2. \end{aligned}$$

∎

Proof of Lemma 2 By Definition (3), for any $\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{z}' \in \mathcal{C}$, $\mathbf{v} \in \partial h(\mathbf{prox}(\mathbf{x}))$, and $\mathbf{w} \in \partial h(\mathbf{prox}(\mathbf{y}))$, we have

$$\begin{aligned} \langle \mathbf{v}, \mathbf{z} - \mathbf{prox}(\mathbf{x}) \rangle &\geq -\langle \nabla f(\mathbf{x}) + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{z} - \mathbf{prox}(\mathbf{x}) \rangle, \\ \langle \mathbf{w}, \mathbf{z}' - \mathbf{prox}(\mathbf{y}) \rangle &\geq -\langle \nabla f(\mathbf{y}) + L(\mathbf{prox}(\mathbf{y}) - \mathbf{y}), \mathbf{z}' - \mathbf{prox}(\mathbf{y}) \rangle. \end{aligned}$$

In particular, for $\mathbf{z} = \mathbf{prox}(\mathbf{y})$ and $\mathbf{z}' = \mathbf{prox}(\mathbf{z})$, we get

$$\begin{aligned} \langle \mathbf{v}, \mathbf{prox}(\mathbf{y}) - \mathbf{prox}(\mathbf{x}) \rangle &\geq -\langle \nabla f(\mathbf{x}) + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{prox}(\mathbf{y}) - \mathbf{prox}(\mathbf{x}) \rangle, \\ \langle \mathbf{w}, \mathbf{prox}(\mathbf{y}) - \mathbf{prox}(\mathbf{x}) \rangle &\leq \langle \nabla f(\mathbf{y}) + L(\mathbf{prox}(\mathbf{y}) - \mathbf{y}), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle. \end{aligned}$$

By monotonicity of sub-gradient, we get

$$\langle \mathbf{v}, \mathbf{prox}(\mathbf{y}) - \mathbf{prox}(\mathbf{x}) \rangle \leq \langle \mathbf{w}, \mathbf{prox}(\mathbf{y}) - \mathbf{prox}(\mathbf{x}) \rangle.$$

So

$$\langle \nabla f(\mathbf{x}) + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle \leq \langle \nabla f(\mathbf{y}) + L(\mathbf{prox}(\mathbf{y}) - \mathbf{y}), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle,$$

and as a result

$$\begin{aligned}
&\langle \nabla f(\mathbf{x}) + L(\mathbf{prox}(\mathbf{x}) - \mathbf{x}), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle \\
&= \langle \nabla f(\mathbf{x}) + L\left(\mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) + \mathbf{prox}(\mathbf{y}) - \mathbf{x}\right), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle \\
&= L\|\mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y})\|_2^2 + \langle \nabla f(\mathbf{x}) + L(\mathbf{prox}(\mathbf{y}) - \mathbf{x}), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle \\
&\leq \langle \nabla f(\mathbf{y}) + L(\mathbf{prox}(\mathbf{y}) - \mathbf{y}), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle,
\end{aligned}$$

which gives

$$\begin{aligned}
L\|\mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y})\|_2^2 &\leq \langle \nabla f(\mathbf{y}) - \nabla f(\mathbf{x}) + L(\mathbf{x} - \mathbf{y}), \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y}) \rangle \\
&\leq \left(\|\nabla f(\mathbf{y}) - \nabla f(\mathbf{x})\|_2 + L\|\mathbf{x} - \mathbf{y}\|_2\right) \|\mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y})\|_2 \\
&\leq 2L\|\mathbf{x} - \mathbf{y}\|_2 \|\mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{y})\|_2,
\end{aligned}$$

and the result follows. ∎

Proof of Lemma 3 Items i and ii, are simply Steps 2 and 3, respectively. For item iii, we have

$$\|\mathbf{x}-\mathbf{w}\|_2 = \|t\mathbf{y}+(1-t)\mathbf{z}-t^*\mathbf{y}-(1-t^*)\mathbf{z}\|_2 = \|(t-t^*)\mathbf{y}-(t-t^*)\mathbf{z}\|_2 = |t-t^*|\|\mathbf{y}-\mathbf{z}\|_2 \leq \epsilon\|\mathbf{y}-\mathbf{z}\|_2.$$

Now it follows that

$$\begin{aligned}
|\langle \mathbf{prox}(\mathbf{x}) - \mathbf{x}, \mathbf{y} - \mathbf{z} \rangle| &= |\langle \mathbf{prox}(\mathbf{x}) - \mathbf{x}, \mathbf{y} - \mathbf{z} \rangle - \langle \mathbf{prox}(\mathbf{w}) - \mathbf{w}, \mathbf{y} - \mathbf{z} \rangle| \\
&\leq \|\langle \mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{w}), \mathbf{y} - \mathbf{z} \rangle\|_2 + |\langle \mathbf{x} - \mathbf{w}, \mathbf{y} - \mathbf{z} \rangle| \\
&\leq \|\mathbf{prox}(\mathbf{x}) - \mathbf{prox}(\mathbf{w})\|_2 \|\mathbf{y} - \mathbf{z}\|_2 + \|\mathbf{x} - \mathbf{w}\|_2 \|\mathbf{y} - \mathbf{z}\|_2 \\
\text{By Lemma 2} \quad &\leq 2\|\mathbf{x} - \mathbf{w}\|_2 \|\mathbf{y} - \mathbf{z}\|_2 + \|\mathbf{x} - \mathbf{w}\|_2 \|\mathbf{y} - \mathbf{z}\|_2 \\
&= 3\|\mathbf{x} - \mathbf{w}\|_2 \|\mathbf{y} - \mathbf{z}\|_2 \\
&\leq 3\epsilon\|\mathbf{y} - \mathbf{z}\|_2^2.
\end{aligned}$$

∎

Proof of Corollary 1 Note that by Step 4 of Algorithm 1), $\mathbf{p}_k = -L(\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k)$. For $k = 1$, since $\mathbf{x}_1 = \mathbf{y}_1 = \mathbf{z}_1$, the inequality is trivially true. For $k \geq 2$, we consider the three cases of Corollary 3: (i) if $\mathbf{x}_k = \mathbf{y}_k$, the right hand side is $1/T \geq 0$ and the left hand side is $\langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k \rangle = \langle -L(\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k), \mathbf{x}_k - \mathbf{z}_k \rangle \leq 0$, (ii) if $\mathbf{x}_k = \mathbf{z}_k$, the left hand side is 0 and

$\langle \mathbf{p}_k, \mathbf{y}_k - \mathbf{x}_k \rangle = \langle -L(\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k), \mathbf{y}_k - \mathbf{x}_k \rangle \geq 0$, so the inequality holds trivially, and (iii) in this last case, for some $t \in (0,1)$, we have

$$\begin{aligned}\langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k \rangle &= \langle -L(\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k), t\mathbf{y}_k + (1-t)\mathbf{z}_k - \mathbf{z}_k \rangle \\ &= -Lt\langle (\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k), \mathbf{y}_k - \mathbf{z}_k \rangle,\end{aligned}$$

and

$$\begin{aligned}\langle \mathbf{p}_k, \mathbf{y}_k - \mathbf{x}_k \rangle &= \langle -L(\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k), \mathbf{y}_k - t\mathbf{y}_k - (1-t)\mathbf{z}_k \rangle \\ &= -L(1-t)\langle (\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k), (\mathbf{y}_k - \mathbf{z}_k) \rangle.\end{aligned}$$

Hence

$$\begin{aligned}&\langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k \rangle - (\eta_k L_k - 1)\langle \mathbf{p}_k, \mathbf{y}_k - \mathbf{x}_k \rangle \\ &\leq |\langle \mathbf{p}_k, \mathbf{x}_k - \mathbf{z}_k \rangle - (\eta_k L_k - 1)\langle \mathbf{p}_k, \mathbf{y}_k - \mathbf{x}_k \rangle| \\ &= |(-Lt + (\eta_k L_k - 1)L(1-t))\langle (\mathbf{prox}(\mathbf{x}_k) - \mathbf{x}_k), (\mathbf{y}_k - \mathbf{z}_k) \rangle| \\ &\leq 3|(-Lt + (\eta_k L_k - 1)L(1-t))|\|\mathbf{y}_k - \mathbf{z}_k\|_2^2 \epsilon_k \\ &= 3|\eta_k L_k (1-t) + 1|L\|\mathbf{y}_k - \mathbf{z}_k\|_2^2 \epsilon_k \\ &= 3(\eta_k L_k + 1)L\|\mathbf{y}_k - \mathbf{z}_k\|_2^2 \epsilon_k \\ &= 6\eta_k L_k L\|\mathbf{y}_k - \mathbf{z}_k\|_2^2 \epsilon_k \\ &= \frac{6D\eta_k L_k L\|\mathbf{y}_k - \mathbf{z}_k\|_2^2}{D} \frac{1}{6dT^3} \\ &\leq \frac{DL\eta_k L_k}{T^3},\end{aligned}$$

where in the last line we used the fact that $\|y_k - z_k\|_2^2 \leq Dd$ ∎

Proof of Lemma 4 For any $\mathbf{u} \in \mathcal{C}$ and by optimality of $\mathbf{z}_{k+1}$, we have

$$\langle \eta_k \mathbf{p}_k, \mathbf{z}_{k+1} - \mathbf{u} \rangle \leq \langle S_k(\mathbf{z}_{k+1} - \mathbf{z}_k), \mathbf{u} - \mathbf{z}_{k+1} \rangle.$$

Hence, using (4) and (5), it follows that

$$\begin{aligned}\langle \eta_k \mathbf{p}_k, \mathbf{z}_k - \mathbf{u} \rangle &= \langle \eta_k \mathbf{p}_k, \mathbf{z}_k - \mathbf{z}_{k+1} \rangle + \langle \eta_k \mathbf{p}_k, \mathbf{z}_{k+1} - \mathbf{u} \rangle \\ &\leq \langle \eta_k \mathbf{p}_k, \mathbf{z}_k - \mathbf{z}_{k+1} \rangle - \langle S_k(\mathbf{z}_{k+1} - \mathbf{z}_k), \mathbf{z}_{k+1} - \mathbf{u} \rangle \\ &= \langle \eta_k \mathbf{p}_k, \mathbf{z}_k - \mathbf{z}_{k+1} \rangle - \frac{1}{2}\|\mathbf{z}_{k+1} - \mathbf{z}_k\|_{S_k}^2 - \frac{1}{2}\|\mathbf{z}_{k+1} - \mathbf{u}\|_{S_k}^2 + \frac{1}{2}\|\mathbf{u} - \mathbf{z}_k\|_{S_k}^2 \\ &\leq \sup_{\mathbf{z} \in \mathbb{R}^d} \left\{ \langle \eta_k \mathbf{p}_k, \mathbf{z} \rangle - \frac{1}{2}\|\mathbf{z}\|_{S_k}^2 \right\} - \frac{1}{2}\|\mathbf{z}_{k+1} - \mathbf{u}\|_{S_k}^2 + \frac{1}{2}\|\mathbf{u} - \mathbf{z}_k\|_{S_k}^2 \\ &= \frac{\eta_k^2}{2}\|\mathbf{p}_k\|_{S_k^*}^2 - \frac{1}{2}\|\mathbf{u} - \mathbf{z}_{k+1}\|_{S_k}^2 + \frac{1}{2}\|\mathbf{u} - \mathbf{z}_k\|_{S_k}^2.\end{aligned}$$

Now recalling from from Steps 6- 8 of Algorithm 1 that $S_k = \text{diag}(\mathbf{s}_k) + \delta\mathbb{I}$ and $\mathbf{s}_k \geq \mathbf{s}_{k-1}$, we sum over k to get

$$\begin{aligned}
\sum_{k=1}^{T}\langle \eta_k \mathbf{p}_k, \mathbf{z}_k - u\rangle &\leq \sum_{k=1}^{T}\frac{\eta_k^2}{2}\|\mathbf{p}_k\|_{S_k^*}^2 + \frac{1}{2}\|\mathbf{u}-\mathbf{z}_1\|_{S_1}^2 + \sum_{k=2}^{T}\frac{1}{2}\|\mathbf{u}-\mathbf{z}_k\|_{S_k}^2 - \frac{1}{2}\|\mathbf{u}-\mathbf{z}_k\|_{S_{k-1}}^2 \\
&= \sum_{k=1}^{T}\frac{\eta_k^2}{2}\|\mathbf{p}_k\|_{S_k^*}^2 + \frac{1}{2}\|\mathbf{u}-\mathbf{z}_1\|_{S_1}^2 + \frac{1}{2}\sum_{k=2}^{T}\langle (S_k - S_{k-1})(\mathbf{u}-\mathbf{z}_k), \mathbf{u}-\mathbf{z}_k\rangle \\
&\leq \sum_{k=1}^{T}\frac{\eta_k^2}{2}\|\mathbf{p}_k\|_{S_k^*}^2 + \frac{1}{2}\|\mathbf{u}-\mathbf{z}_1\|_{\infty}^2\langle \mathbf{s}_1, \mathbf{1}\rangle + \frac{1}{2}\sum_{k=2}^{T}\|\mathbf{u}-\mathbf{z}_k\|_{\infty}^2\langle \mathbf{s}_k - \mathbf{s}_{k-1}, \mathbf{1}\rangle \\
&\leq \sum_{k=1}^{T}\frac{\eta_k^2}{2}\|\mathbf{p}_k\|_{S_k^*}^2 + \frac{D}{2}\langle s_1, \mathbf{1}\rangle + \frac{D}{2}\sum_{k=2}^{T}\langle \mathbf{s}_k - \mathbf{s}_{k-1}, \mathbf{1}\rangle \\
&= \sum_{k=1}^{T}\frac{\eta_k^2}{2}\|\mathbf{p}_k\|_{S_k^*}^2 + \frac{D}{2}\|\mathbf{s}_T\|_1
\end{aligned}$$

∎

Proof of Lemma 5 To prove part i, we use the following inequality introduced in the proof of Lemma 4 in [7]: for any arbitrary real-valued sequence of $\{a_i\}_{i=1}^T$ and its vector representation as $a_{1:T} = [a_1, a_2, \ldots, a_T]$, we have

$$\sum_{k=1}^{T}\frac{a_k^2}{\|a_{1:k}\|_2} \leq 2\|a_{1:T}\|_2.$$

So it follows that

$$\sum_{k=1}^{T}\mathbf{g}_k^T S_k^{-1}\mathbf{g}_k = \sum_{k=1}^{T}\sum_{i=1}^{d}\frac{\mathbf{g}_k^2(i)}{\mathbf{s}_k^2(i)} = \sum_{i=1}^{d}\sum_{k=1}^{T}\frac{\mathbf{g}_k^2(i)}{\mathbf{s}_k(i)} = \sum_{i=1}^{d}\sum_{k=1}^{T}\frac{\mathbf{g}_k^2(i)}{\|G_k(i,:)\|_2} \leq 2q_T,$$

where the last equality follows from the definition of $\mathbf{s}_k$ in Step 7 of Algorithm 1.

For the rest of the proof, one can easily see that

$$\sum_{k=1}^{T}\mathbf{g}_k^T S^{-1}\mathbf{g}_k = \sum_{k=1}^{T}\sum_{i=1}^{d}\frac{\mathbf{g}_k^2(i)}{\mathbf{s}(i)} = \sum_{i=1}^{d}\frac{a(i)}{\mathbf{s}(i)},$$

where $a(i) := \sum_{k=1}^{T}\mathbf{g}_k^2(i)$ and $\mathbf{s} = \text{diag}(S)$. Now the Lagrangian for $\lambda \geq 0$ and $\boldsymbol{\nu} \geq \mathbf{0}$, can be written as

$$\mathcal{L}(\mathbf{s}, \lambda, \boldsymbol{\nu}) = \sum_{i=1}^{d}\frac{a(i)}{\mathbf{s}(i)} + \lambda\left(\sum_{i=1}^{d}\mathbf{s}(i) - 1\right) + \langle \boldsymbol{\nu}, \mathbf{s}\rangle.$$

Since the strong duality holds, for any primal-dual optimal solutions, S^*, λ^* and $\boldsymbol{\nu}^*$, it follows from complementary slackness that $\boldsymbol{\nu}^* = \mathbf{0}$ (since $\mathbf{s}^* > \mathbf{0}$). Now requiring that $\partial \mathcal{L}(\mathbf{s}^*, \lambda^*, \boldsymbol{\nu}^*)/\partial \mathbf{s}(i) = 0$ gives $\lambda^* \mathbf{s}^*(i) = \sqrt{a_i} > 0$, which since $\mathbf{s}^*(i) > 0$, implies that $\lambda^* > 0$. As a result, by using complementary slackness again, we must have $\sum_{i=1}^d \mathbf{s}^*(i) = 1$. Now simple algebraic calculations gives $\mathbf{s}^*(i) = \sqrt{a_i}/(\sum_{i=1}^d \sqrt{a_i})$ and part ii follows.

For part iii, recall that $\|\mathbf{g}_k\|_2 = 1$. Now, since $\lambda_{\min}(S^{01}) \geq 1$, one has $1 \leq \mathbf{g}_k^T S^{-1} \mathbf{g}_k$, and so $q_{_T} \geq 1$. One the other hand, consider the optimization problem

$$\begin{aligned} \max \sum_{i=1}^d \|G_T(i,:)\|_2 &= \sum_{i=1}^d \sqrt{\sum_{k=1}^T \mathbf{g}_i^2(k)} \\ \text{s.t. } \|\mathbf{g}_k\|_2^2 &= 1, \ k = 1, 2, \ldots, T. \end{aligned}$$

The Lagrangian can be written as

$$\mathcal{L}(\{\mathbf{g}_k\}_{k=1}^T, \{\lambda\}_{k=1}^T) = \sum_{i=1}^d \sqrt{\sum_{k=1}^T \mathbf{g}_i^2(k)} + \sum_{k=1}^T \lambda_k \left(1 - \sum_{i=1}^d \mathbf{g}_i^2(k)\right).$$

By KKT necessary condition, we require that $\partial \mathcal{L}(\{\mathbf{g}_k\}_{k=1}^T, \{\lambda\}_{k=1}^T)/\partial \mathbf{g}_i(k) = 0$, which implies that

$$\lambda_k = \frac{1}{2\sqrt{\sum_{k=1}^T \mathbf{g}_i^2(k)}}, \quad i = 1, 2, \ldots, d.$$

Hence,

$$T = \sum_{i=1}^d \sum_{k=1}^T \mathbf{g}_i^2(k) = \frac{d}{4\lambda_k^2},$$

and so $2\lambda_k = \sqrt{d/T}$, which gives $q_{_T} \leq \sqrt{dT}$. ∎

Proof of Lemma 7 We prove by induction on T. For $T = 1$, we have $\eta_1 = 1/L_1$, and the base case holds trivially. Suppose the desired relation holds for $T - 1$. We have

$$\begin{aligned} \sum_{k=1}^T \eta_k &= \sum_{k=1}^{T-1} \eta_k + \eta_T \\ \text{(By induction hypothesis)} \ &\geq \frac{(T-1)^3}{1000 \sum_{k=1}^{T-1} L_k} + \frac{1}{2L_T} + \sqrt{\frac{1}{4L_T^2} + \frac{(T-1)^3}{1000 L_T \sum_{k=1}^{T-1} L_k}} \\ &\geq \frac{(T-1)^3}{1000 \sum_{k=1}^{T-1} L_k} + \sqrt{\frac{(T-1)^3}{1000 L_T \sum_{k=1}^{T-1} L_k}} \\ &\geq \frac{(T-1)^3}{1000 \sum_{k=1}^{T-1} L_k} + \sqrt{\frac{T^3}{8000 L_T \sum_{k=1}^{T-1} L_k}}. \end{aligned}$$

Now if

$$\frac{(T-1)^3}{1000\sum_{k=1}^{T-1} L_k} \geq \frac{T^3}{1000\sum_{k=1}^{T} L_k},$$

then we are done. Otherwise denoting $\alpha := \sum_{k=1}^{T} L_k$, we must have that

$$\begin{aligned} L_T &\leq \frac{\alpha T^3 - \alpha(T-1)^3}{T^3} \\ &= \frac{\alpha T^3 - \alpha\left(T^3 - 3T^2 + 3T - 1\right)}{T^3} \\ &= \frac{\alpha(3T^2 - 3T + 1)}{T^3} \\ &\leq \frac{4\sum_{k=1}^{T} L_k}{T}. \end{aligned}$$

Hence, we get

$$\begin{aligned} \sum_{k=1}^{T} \eta_k &\geq \frac{(T-1)^3}{1000\sum_{k=1}^{T-1} L_k} + \sqrt{\frac{T^4}{32000 L_T \left(\sum_{k=1}^{T} L_k\right)^2}} \\ &\geq \frac{(T-1)^3}{1000\sum_{k=1}^{T} L_k} + \frac{4T^2}{1000\sum_{k=1}^{T} L_k} \\ &\geq \frac{T^3}{1000\sum_{k=1}^{T} L_k}. \end{aligned}$$

∎